

# ANALYZING TECHNICAL REQUIREMENTS

**After reading this chapter and completing the exercises, you will be able to:**

- ◆ Analyze an organization's existing and planned technical environment and goals
- ◆ Analyze the impact of infrastructure design on the existing and planned technical environment
- ◆ Analyze network requirements for client computer access
- ◆ Analyze the existing disaster recovery strategy for client computers, servers, and the network

In Chapter 2 you learned how to analyze the business requirements of an organization. The business analysis information was gathered primarily from nontechnical managers and end users. In looking at the existing and planned business models, company processes, and organizational structures, you were trying to determine who is communicating with whom and how. In Chapter 3, you continue in that vein in your technical analysis as you look at user and resource distribution in an organization. In this chapter, you are concerned with network access patterns and how a network design can improve on existing access patterns. The information you require for a technical analysis is gathered primarily from technical staff.

The objectives for this chapter map to exam objectives. In a real-world scenario, performing each analysis as described in the objectives would be an extremely ambitious undertaking, perhaps requiring a large, full-time, highly paid staff and/or large teams of consultants. Obviously, you are only one person. Thus, we restrict this chapter to the information that would help you understand *how to begin evaluating* an organization's technical requirements, limiting our scope to what truly impacts the planned Windows 2000 environment.

---

## ANALYZE CURRENT AND PLANNED TECHNICAL ENVIRONMENT AND GOALS

Like most analysis, technical analysis for a network design starts with learning the current status. This is a logical start—you can't plan where you're going unless you know where you are. This rule applies whether you are looking at the map of a city or a map of a corporate network infrastructure. To understand the current technical environment, you must analyze the following areas:

- Company size in terms of end-user and resource distribution
- Available connectivity between the geographic location of work sites and remote sites
- Net available bandwidth and latency issues
- Performance, availability, reliability, and scalability requirements of services
- Data and system access patterns
- Network roles and responsibilities
- Security considerations

### Analyze Company Size in Terms of End-User and Resource Distribution

Although Microsoft lists company size as a significant factor, the size of the company, measured solely in the number of employees, is not very significant because not all employees have jobs that require access to computers. What is significant is the number of employees with computer access and their distribution on the network relative to the location of the network resources they use. We will refer to these employees as end users.

Network resources accessed by end users are applications, files, printers, and databases accessed *over the network*. They specifically exclude the applications, files, printers, and databases that reside on each computer and that are not shared on the network. As you look for these resources and ask people for their information, remember that people may forget to mention applications, printers, and other resources used solely by small groups of users. Make sure you ferret out these “hidden resources.” Their existence, location, and how they are accessed have a direct impact on the type and volume of network traffic for those groups of users.

You may find this information about end users and network resources hard to obtain at times, believe it or not. For instance, if a network has grown faster than the ability of technical staff to support and document it, real detective skills may be required. This can delay the entire process and be costly. As you work through this process, impress on managers the importance of obtaining the most accurate information in order to keep costs down and have a successful network upgrade.

Human roadblocks to the information-gathering process can come in many forms. One is the “de facto” support person in a department who loves his or her unofficial technical support job more than his or her real job, and who is afraid that if IT becomes more formal and more successful, they will no longer have a paid hobby. Another possible human roadblock is the department manager who has been “creative” in finding the money to add resources and clients to the network, and would rather not have a formal accounting of their network usage. There are many other personnel issues that can trip you up. Watch for them.

Some useful questions at this stage of the information-gathering process include the following:

- What applications are in use?
- What files are accessed with high frequency?
- Where are the network printers?
- Where are the databases for this network?
- Are desktop computers sharing file and printer resources?
- Are most users concentrated in just a few locations?
- Are users located at a large number of physical locations?
- Where are the network resources relative to the user locations?
- Are there identifiable availability problems?
- What growth is expected in the number of users, locations of users, and resources?
- What changes are planned in resource usage and applications usage?

You will want to organize the information you gather from these questions. You may use simple forms as you gather the information. Then you will want to create graphical representation of the network. Eventually a complete picture of the network will come into focus.

Table 3-1 is an example of a form with information gathered about a fictitious novelty/gift company, NovlGifts. They have two separate markets for their products—a line of collectibles sold to consumers through infomercials and newspaper supplement ads, and another line of collectibles sold at wholesale to gift shops in the eastern United States. Their corporate office and fulfillment center is in Boston, MA, with regional offices in New York, NY, Winston-Salem, NC, and Miami, FL. They have more than 1000 end users, each with their own computers (“Hosts” in Table 3-1) and distributed as indicated. They presently have all sites connected to Boston through leased lines. The Boston site is connected to the Internet through a T1 line to the provider’s site. They currently have an NT domain that will be migrated to Windows 2000 Active Directory. Some servers are providing multiple roles, so there are fewer physical servers than indicated under “Number and Types of Servers.”

**Table 3-1** Current user and resource distribution for NovlGifts

Location	Current Number of Hosts	Existing Connections to Boston (type and speed)	Number and Types of Servers
HQ – Boston	825	N/A	File and print servers: 6 Mail servers: 1 Domain controllers: 2* WINS servers: 1** Databases (sales order/entry and inventory): 1** Database (accounting/finance): 1** DNS: 2**
New York	75	128 Kbps	File and print servers: 1*** Domain controller: 1***
Winston-Salem	81	128 Kbps	File and print servers: 1*** Domain controller: 1***
Miami	71	128 Kbps	File and print servers: 1*** Domain controller: 1***

\* At the headquarters, there are two Windows NT domain controllers (the PDC and one BDC), each of which is a WINS server.

\*\* These three applications/services are hosted on three separate Unix servers.

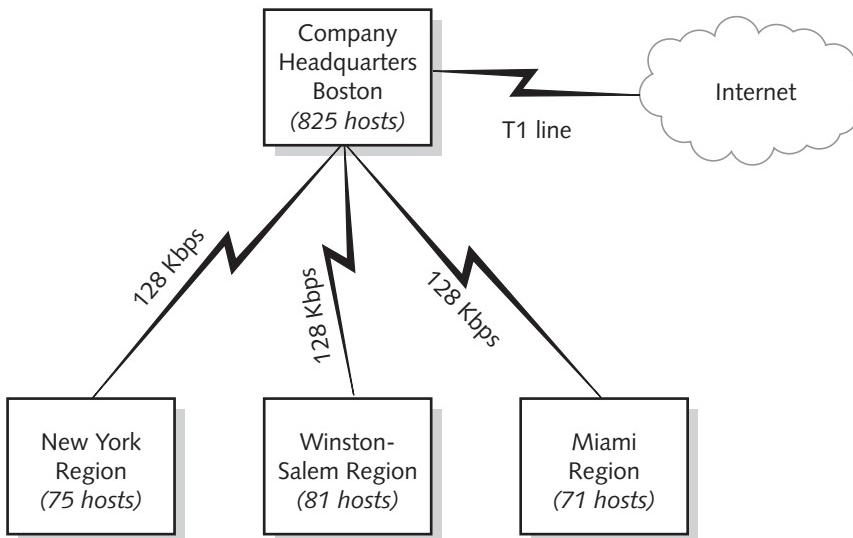
\*\*\* Each regional office has a single NT server that serves both roles.

## Analyze Connectivity Between Geographic Locations

Although you may know that the New York office is connected to the Boston office by a 128 Kbps leased line, you also need to know how reliable the connection is, the cost and present usage of the connection, and if there is room for additional traffic.

When you are looking at whether there is room for additional traffic, you will find that sometimes the options are limited because the providers of the WAN services may only have so many choices for connections, especially in remote areas throughout the world. Limitations such as these must be incorporated into your analysis.

Consider Figure 3-1. In it, you see that all the sites of NovlGifts are connected through leased lines to the corporate headquarters. This network was designed to support the model of primarily client/server applications. At the time of its implementation in 1995, only headquarters personnel used the Internet. Now users at all sites access the Internet through the headquarter's Internet connection. Users now are complaining of slow access, and tests have shown that at peak usage times, all the connections are overused.



**Figure 3-1** Present network of NovlGifts

The following are questions that you might ask IT staff (or yourself!) as you dig out information related to geographic locations and your network:

- What type of WAN service do you have?
- Are your users and managers happy with the performance of the WAN connections?
- Will new requirements involve WAN traffic?

The information you gather in this stage will be used in the next analysis, the assessment of net available bandwidth and latency issues. Collectively, this information will become the basis for the Resource leg of the Iron Triangle.

## Analyze Net Available Bandwidth and Latency Issues

Let's begin this discussion by defining bandwidth and latency. **Bandwidth** is the amount of data that can be transmitted in a fixed amount of time. Bandwidth and capacity are often used synonymously when talking about networks. **Latency** in a network is the amount of time it takes data to travel from source to destination. When you know both pieces of information—bandwidth and latency—you have a pretty clear picture of the network.

### Net Available Bandwidth

In digital network connections, the bandwidth is usually measured in millions of bits per second (Mbps). For example, a T1 connection has a bandwidth of 1.544 Mbps. OK, we admit that simply knowing that you have a 1.544 Mbps T1 connection to a phone company's WAN and 128 Kbps connections from your branch offices to the WAN is not

enough. You must also determine the actual bandwidth of the T1 connection because the design team needs current, valid figures with which to work.

Determining “actual bandwidth” can be a little tricky. Many factors can make your available bandwidth much less than the line speed. One factor is line service from your WAN provider. The service could be Frame Relay, SONET, or ATM. Each of these services has characteristics that affect the quality of the bandwidth. Therefore, the choice of WAN services offered by your provider can have a significant impact on performance.

Let’s examine a “for instance” that illustrates the variability. Say you have a 1.544 Mbps ATM circuit; ATM always uses fixed-length packets called “cells”. These cells are 53 bytes long and have a 5-byte header. That means that about 10% of your total bandwidth is used by the ATM protocol itself. When you add your IP header, TCP header, and SMB header on top of that, a simple file-copy operation from an NT server to your PC can actually use more than twice as much bandwidth as the size of the file.

When evaluating WAN information, remember that some layer 2 WAN protocols offer best-effort service, while others offer guaranteed service, which means that some are more efficient because of lower overhead—until packets start dropping, of course. On a very unstable, low-bandwidth link, a protocol that guarantees delivery by handling error detection and retransmission at layer 2, such as the X.25 protocol, may be the better, if less obvious, choice.



As a customer of the WAN provider, your perspective may be that everything beyond the T1 is an unknown and that the WAN is a cloud for that reason. However, service providers often can see beyond the T1, so they don't generally refer to it as a cloud.

Once you know all the caveats, it’s time to find out the actual usable bandwidth. There are many tools for doing so. You can use Network Associates’ Sniffer ([www.sniffer.com](http://www.sniffer.com)) products for all network technologies and 3Com’s ([www.3com.com](http://www.3com.com)) Transcend products if your organization uses 3Com network devices. For a hardware-based protocol analyzer, look at Wandel & Golterman’s Domino product line. (Thanks to mergers, acquisitions, and name changes, W&G products can now be found under Acterna at [www.acterna.com/products/domino/domino\\_lan.html](http://www.acterna.com/products/domino/domino_lan.html).) Last but not least, don’t forget that Microsoft’s Network Monitor and System Monitor can give you some network performance information.



Try Hands-on Project 3-3 and Hands-on Project 3-6 for practice on monitoring.

You may wish to supplement your analysis with information from your WAN provider. For instance, the provider of your WAN connection may also provide reports that will show circuit loss information about the quality of the connection on the monthly statement. The provider may also have additional tools, often available through a Web site, that you can use to initiate monitoring of the circuit. Although some network managers express

distrust of the reliability of these sources, others do give them some credence, especially if they do not have the resources to conduct their own tests.

Remember that packet loss varies exponentially with bandwidth. At 10% network utilization, you may have no packet loss. At 50%, you may lose 10 packets per million, at 80% you may lose a quarter of your total packets, and at 90% you may lose half your total packets. So obviously, if your link is currently 20% utilized and you have a loss of 50 packets per million, you cannot extrapolate that accurately to get an available bandwidth number.

You can, however, use this information to determine your effective bandwidth now and your bandwidth cushion for future use. See Figure 3-2 for the formula.

$$\text{Theoretical maximum bandwidth} - \text{circuit loss} = \text{gross available bandwidth}$$

$$\text{gross available bandwidth} - \text{other existing traffic (applications and services)} = \text{net available bandwidth}$$

$$\text{net available bandwidth} - \text{required new applications bandwidth} = \text{bandwidth cushion available for future needs}$$

**Figure 3-2** Bandwidth formula

Let's put the formula to use. For example, assume we have a 1.544 Mbps point-to-point connection and we have discovered an average circuit loss of 3%, and that our existing traffic has an average load of .81 Mbps. In addition, we have calculated that the bandwidth for the required new applications will require .42 Mbps additional traffic. Figure 3-3 illustrates the calculation based on this information.

$$1.544 \text{ Mbps} - .05 \text{ Mbps} = 1.494 \text{ Mbps gross available bandwidth}$$

$$1.494 \text{ Mbps} - .81 \text{ Mbps} = .684 \text{ Mbps net available bandwidth}$$

$$.684 \text{ Mbps} - .42 \text{ Mbps} = .264 \text{ Mbps bandwidth cushion}$$

**Figure 3-3** Bandwidth calculation

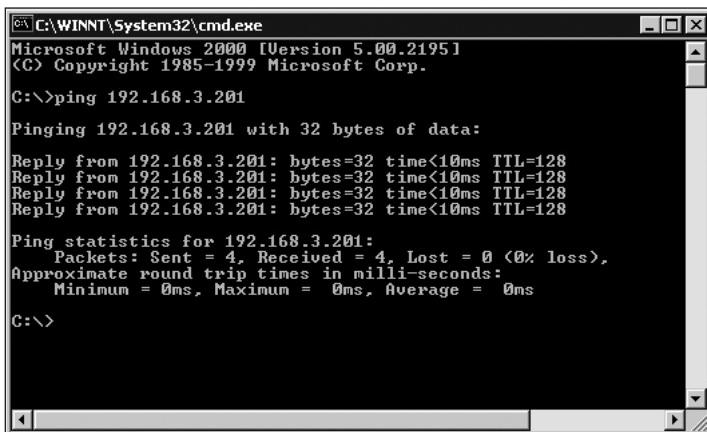
Although this is still just a ballpark figure and would require further research to discover peak usage rates, we can see that the .264 Mbps bandwidth cushion, which is based on average usage, would be adequate during peak usage times.

## Latency Issues

Networks should have a minimum latency, and the minimum latency should not vary significantly. Latency variation is referred to as "jitter" and can be more detrimental to some applications than latency alone. Multimedia applications are an example of a category of fragile applications that cannot withstand jitter.

The path your traffic must take directly impacts the success of the applications. As an example, satellite links have the greatest latency, because of the time required for a signal to travel from a transmitter on the ground to the satellite, and then from the satellite to the destination receiver. On the other hand, when we are using simple client/server applications, we are not usually affected by latency. These applications are often simply moving files across the network, using protocols with long time-out values and multiple retry periods.

In Figure 3-4, we used the ping command to test latency across a local router. In this case the latency is less than 10 milliseconds (ms). In Figure 3-5, we used the ping command to test latency to a host on the Internet. Notice the increased latency in the results from the Internet test. Put simply, the more routers, and the busier the routers (full buffers), the greater the latency.



```
C:\WINNT\System32\cmd.exe
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-1999 Microsoft Corp.

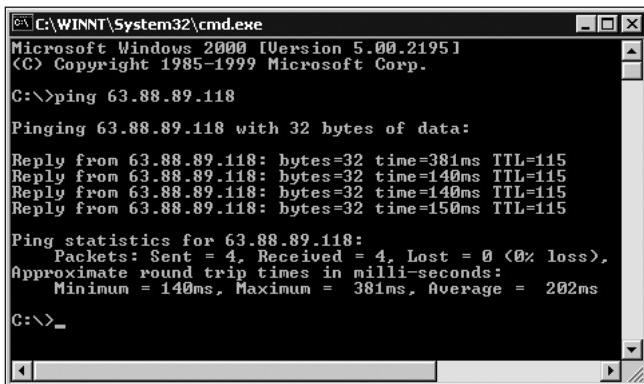
C:>ping 192.168.3.201

Pinging 192.168.3.201 with 32 bytes of data:
Reply from 192.168.3.201: bytes=32 time<10ms TTL=128

Ping statistics for 192.168.3.201:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:>
```

Figure 3-4 Testing for latency across a local router



```
C:\WINNT\System32\cmd.exe
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-1999 Microsoft Corp.

C:>ping 63.88.89.118

Pinging 63.88.89.118 with 32 bytes of data:
Reply from 63.88.89.118: bytes=32 time=381ms TTL=115
Reply from 63.88.89.118: bytes=32 time=140ms TTL=115
Reply from 63.88.89.118: bytes=32 time=140ms TTL=115
Reply from 63.88.89.118: bytes=32 time=150ms TTL=115

Ping statistics for 63.88.89.118:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 140ms, Maximum = 381ms, Average = 202ms

C:>
```

Figure 3-5 Testing for latency to a host on the Internet



See Hands-on Project 3-2 for practice on working with latency.

Other factors that contribute to latency include encryption, encapsulation, compression, segmentation and reassembly of the data, queuing, and large MTUs. However, true latency in the corporate network environment is usually so small that it only affects real-time applications, like multimedia applications. If your network must support such applications, then take a closer look at latency.



Don't confuse latency with user-perceived latency, which means users are reporting that the network is slow, when the real cause may be found in the client and server computers themselves. These are harder to track down, but can include such causes as poorly written device drivers and poorly configured systems (such as too many protocols). Other causes for user-perceived latency may be contention on the system bus, insufficient memory (requiring paging to disk), and nonoptimal disk configuration. Solutions for these problems are beyond the scope of this book.

## Analyze Requirements for Services

The objectives for Microsoft Exam 70-221 list the need to analyze performance, availability, and scalability of services. We have chosen to add functionality, security, and cost to this list to make it more complete and more relevant for the network administrator who is operating in a real-world environment. All are discussed in the following list.

- Business managers and IS staff may ask for performance in terms of bandwidth that should be available for applications and services or in terms of specific technology. You need to understand the ramifications of specific requests. For instance, if they are requesting the use of Voice over IP (VOIP), consider that the ISO has recommended that latency over 150 ms will result in unacceptable voice quality. Although some manufacturers differ, you could end up with an IP telephony project with a very specific requirement of <150 ms delay between a private branch exchange (PBX) and a remote office.
- Availability is the percentage of time the network services are accessible to users. This may be expressed as a percentage of “up” time, say 98.4%, or as a mean time between failures (MTBF) of, say, 4000 hours. Availability also can be expressed in mean time to repair (MTTR). If your network exceeds the MTTR, it is considered to be unavailable.
- Scalability is the ability of the design to shrink or grow with varying levels of demand placed on it. The protocols in use will affect scalability. For instance, NetBEUI is a nice, fast protocol for a small workgroup on a single network segment, but it does not scale well. Since it's not a routable protocol, you can't expand to multiple subnets, as you can with IPX/SPX or TCP/IP.

Name resolution schemes also affect scalability. WINS on NT only scales well to a couple of thousand users. LoadRunner from Mercury Interactive ([www-heva.mercuryinteractive.com](http://www-heva.mercuryinteractive.com)) can help with scalability testing.

- Functionality is the degree to which the network design provides needed results. If one of the desired results was to provide name resolution for NetBIOS clients in a multiple subnetted network, a single WINS server might just provide the functionality. However, when you consider the other requirements of availability, scalability, and performance, you may find that one WINS server simply may not do the trick.
- Security levels need to be defined and agreed upon by the interested parties, and the appropriate security and configuration options must be chosen such that the required level of security is achieved. Security authentication methods might add to network traffic.
- Cost considerations enter the process because of budget restraints, and appropriate decisions about the other requirements must be made. We will not directly deal with cost considerations in this book, because costs for technology are so volatile.

Your interviews with management during your business analysis (see Chapter 2) should have given you the functionality, availability, scalability, security, cost, and performance requirements of the business. We repeat them here to show their technical aspects.

At this point, let's examine the business and technical requirements for our fictitious company, NovlGifts. The business goal are:

- Provide a Web site that gift store customers can use to place new orders and check on the status of pending orders.
- Provide a Web site for the gift line currently sold through TV infomercials and newspaper ads.
- Increase profits by improving the internal processes. This will include a move from an Oracle-based inventory system hosted on a Unix server to a distributed inventory system hosted on a SQL server.
- Increase sales by 35% in the next two years.
- Reduce storage and shipping charges. The strategy to accomplish this is to move away from the current practice of warehousing huge amounts of inventory. To this end, they are negotiating with their suppliers for direct shipments of wholesale orders to their customers.
- Improve the performance of the existing WAN in order to support the new usage.

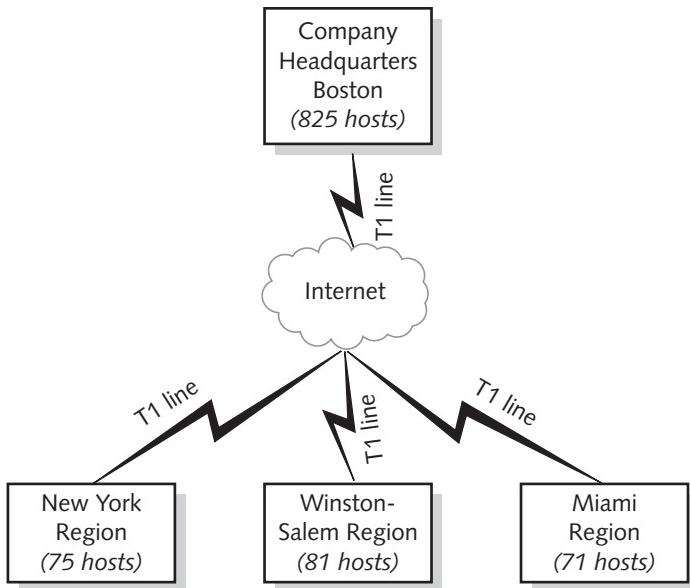
- Ensure the security of data between sites.
- Minimize the cost of the enhanced WAN to the business.

Now, we look at the related technical goals:

- Update the capacity of the WAN to provide improvements in performance, availability, and scalability.
- Provide a Web site to support the new Web-based ordering that is available 99.99% of the time.
- Provide network availability on the internal network, including between sites, to 99.98%.
- Provide a network that is reliable and offers an MTBF of 4500 hours, and an MTTR of two hours.
- Provide a network that will scale to accommodate higher bandwidth in the future.
- Provide a network that will support logical growth in the allocation of IP addresses (this is a scalability requirement).
- Provide a network that will be easier to manage.
- Provide a network that will provide a response time of one-tenth of a second or less to the order-entry system for orders placed internally.
- Prevent unauthorized access to company data that travels over the Internet between sites.
- Prevent users access to untrusted Internet sites.

Once you know the requirements for the business, you measure to determine if you have indeed reached the required levels.

Let's create an example of measuring to determine adherence to a requirement. Let's assume that NovlGifts has removed the leased lines from Boston to the regional offices and has given each office a partial T-1 Internet connection over which they run a virtual private network (VPN). This move means that they meet the requirements for improving the performance of the existing WAN and preventing unauthorized access to company data that travel over the Internet between sites. Figure 3-6 illustrates this change.



**Figure 3-6** VPN of NovlGifts



A tool for evaluating network performance is RoboMon by Heroix ([www.robomon.com](http://www.robomon.com)). This infrastructure management software detects and even corrects complex application, system, and network problems.



The values that you generated in this section go into the Load leg bucket of the Iron Triangle.

## Analyze Data and System Access Patterns

Your analysis of data and system access patterns will show you where and when the network is stressed. You can use Network Monitor, which comes with Windows 2000 (but is not installed by default) or one of several third-party tools. We will mention just a few here:

- LoadRunner from Mercury Interactive comes with several real-time performance monitors, including Transaction Monitor, Server Monitor, Network Delay Monitor, and SNMP Monitor. Transaction Monitor gathers information on average user transaction response time and transaction throughput. Server Monitor shows you where server-specific performance problems may exist on such servers as Web servers, application servers, and database servers. Network Delay Monitor provides a breakdown of network performance by segment. Last, with SNMP Monitor, you can look for performance problems on any SNMP-compliant network component, such as bridges and routers. (You don't want to overlook these devices in your analysis.)

- 3Com has free software tools for use with their equipment, but they also sell DynamicAccess Network Performance Manager & LAN Agent, which does not depend on 3Com products, for real-time analysis of network and server usage and application response times.
- Microsoft's Network Monitor will allow you to capture and analyze network traffic, looking for the types of traffic generated over selected periods of time. Figure 3-7 shows Microsoft Network Monitor displaying summary information after a network capture was performed. In the right pane of the window, you can see the total frames captured, the number of broadcast and multicast frames, the number of frames dropped, and other information. The graph on the upper-left gives you a quick visual summary of traffic characteristics, while the remaining panes provide session and station statistics. You can also use Network Monitor to examine the contents of individual packets, as shown in Figure 3-8, which shows the Network Monitor frame viewer window.

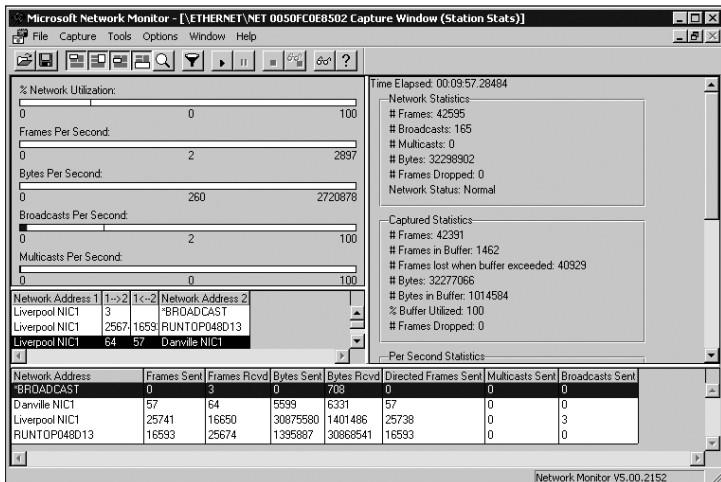
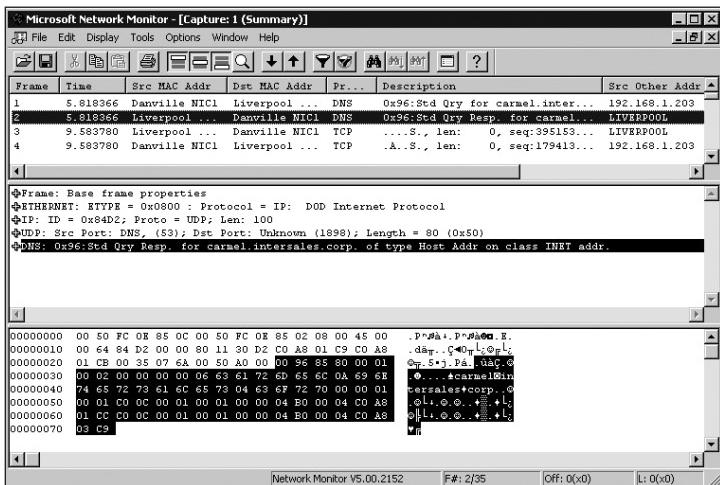


Figure 3-7 Network Monitor capture window

## Analyze Network Roles and Responsibilities

Within your network, Active Directory, while depending on several interrelated services, defines certain roles, especially for Windows 2000 domain controllers. These additional domain controllers are not just understudies. They are equals, more or less. They are equals in that each has a full copy, or replica, of the Active Directory for that domain. The “more or less” comes into play when a domain controller is taking on a special limited role, which can only be played by one Windows 2000 domain controller in a forest or domain.



**Figure 3-8** Network Monitor frame viewer window



The roles in the network should be familiar to anyone reading this book. Thus, only the special roles for Windows 2000 domain controllers are reviewed in detail.

A **global catalog server** is an Active Directory (Windows 2000) domain controller that contains a partial replica of every domain directory partition in the forest as well as a full replica of its own domain directory partition and the schema and configuration directory partitions. This global catalog contains a replica of every object in Active Directory, but only a subset of the attributes of each object. In a multiple domain forest with multiple sites, a global catalog server is required for domain logon by anyone but domain administrators. Therefore, in a multiple domain forest, at least one global catalog server must reside in each site.

Global catalog servers also are used for searches of Active Directory in which you need to search on one or more objects without specifying the domain in which the object resides. Only one global catalog server is created by default; it is created on the first domain controller in the first domain in the forest. However, administrators can assign this role to as many domain controllers in a forest as needed, based on the logon and search requirements mentioned above.

There are actually several **single master operation roles** for Active Directory domain controllers. Two of those roles—schema master and the domain-naming master—can reside only with one domain controller in an entire forest.



The original term used for single master operations was **floating single master operations (FSMOS)**. This term is still used in documentation and in tools that let you move the roles, such as the utility NTDSUTIL.

Although every domain controller in the forest has a copy of the schema for Active Directory, only the **schema master** has a writable copy. By default, the first domain controller in the first domain in the forest has this role. This role can be transferred to another domain controller in the forest as needed. It can be seized with NTDSUTIL, or a similar tool, if the domain controller hosting this role has failed and will never come up on the network again. This role is important because, although changes to the schema are very infrequent, when it is necessary to make a change, the software that allows you to modify the schema focuses on the schema master because that is the only copy of the schema that can be modified.

The **domain-naming master** role is automatically assigned to the first domain controller in the first domain in the forest. The domain-naming master manages the addition or removal of domains in the forest. When you promote a Windows 2000 server to domain controller and make it the first domain controller in a new domain in an existing forest, the Domain Controller Promotion Wizard (DCPROMO) must be able to contact the domain-naming master to have the new domain added to the forest. This role also can be transferred to another domain controller.

Next, you need to consider three more single master operations roles, each of which can exist on one domain controller in each domain in a forest. These roles are the RID master, the PDC emulator, and the infrastructure master.

- The **RID master** is in charge of allocating new relative IDs that are used together with the domain security ID (SID) to create unique security IDs for each object that can be a security principal. Security principals can be user, group, and computer objects.
- The **PDC emulator** has special roles in both mixed-mode and native-mode domains. In a mixed-mode domain, the PDC emulator “pretends” to be a Windows NT 4.0 PDC for the benefit of Windows NT 4.0 BDCs by replicating directory changes to the BDCs in the domain. In both mixed mode and native mode, the PDC emulator servicing requests from pre-Windows 2000 clients would send the information directly to a PDC. This includes any request to modify the directory, such as password changes.

In a native-mode domain, the PDC emulator receives preferential replication of password changes from other Windows 2000 domain controllers (that is the only kind you have in native mode). That way, if a user password is changed and the user attempts to log on before replication of the change has reached the domain controller that is servicing the logon, the authenticating domain controller will not immediately refuse the logon. Instead, it will query the PDC emulator to see if it has recently received a password change for that user.

- The **infrastructure master** is responsible for keeping track of updates of group-to-user references, such as the renaming of a user account when group memberships are changed in different domains. The infrastructure master in the group’s domain registers the updates and replicates them to other infrastructure masters.

There are other roles that are more closely tied to individual services within your network. These are referred to as either roles or services and include the following:

- Name servers include those providing DNS or WINS services. The services they provide include name registration and response to client queries for name IP address resolution. In a Windows 2000 network, name resolution is critical for clients so that they can locate services on the network. Clients use name resolution to find domain controllers for logon, authentication services, and global catalog searches.
- Address assignment servers, or Dynamic Host Configuration Protocol (DHCP) servers, assign IP addresses to client computers and register client host names with Dynamic Name Service (DNS) servers. Placement of these servers will be defined when you plan for IP address allocation on the network. At this point in your analysis, you will determine the location of DHCP servers and the present scheme for address allocation.
- Remote access servers can include servers providing dial-in RAS, VPN, RRAS, or terminal services for remote users. Be sure to determine what types of connections are in use. The connections can include Internet, PSTN, ISDN, DSL, or leased.
- File and print servers are those hosting user home directories, printers, and other data repositories. These are traditionally on the same LAN as the users who access them, per the old 80/20 rule of having 80% of your network traffic within the same LAN, and no more than 20% beyond the LAN. Do not be surprised if you find users accessing servers beyond their LAN for file and print services, especially mobile users. When you find this, be sure to question the need for this traffic. It may be possible to streamline things a bit in this area before adding more services to the network.
- Application servers include database servers, electronic messaging servers, Web servers, and any server hosting the server side of a client/server application.
- Security servers include certificate authority servers, which are used for public key infrastructure authentication of nondomain users, and remote authentication dial-in user service, which is used for authenticating dial-in users.

If there is significant traffic focused on a few servers, then those servers should be placed on high-speed segments.

### Determining the Network Services That Need to Be Added

Using the information gathered in both the business analysis and the technical analysis thus far, you need to determine what services must be provided. If we were executing this step for our fictitious company, NovlGifts, our form might look like Table 3-2.

**Table 3-2** Services that need to be added at NovlGifts

Location	Planned Total Number of Hosts	Planned Connections to the Internet (type/speed)	Number and Types of Servers Planned	Timing of Growth
HQ – Boston	925	T-1 to provider	Web servers: 4 Mail servers: 2 Domain controllers: 2* Global catalog server: 2 WINS servers: 2* DNS servers: 2* Databases (sales order/entry and inventory): 1** Databases (accounting/financial): 1** File and print servers: 7***	Existing
New York	84	Partial T-1 to provider	File and print servers: 1**** Domain controllers: 1****	6 weeks
Winston-Salem	93	Partial T-1 to provider	File and print servers: 1*** Domain controllers: 1****	2 months
Miami	79	Partial T-1 to provider	File and print servers: 1**** Domain controllers: 1****	2 months

\* There are two Windows 2000 domain controllers, each of which is a DNS server. One of the domain controllers has the role of global catalog server, but it is not significant, since this is a single domain forest. This server also has all the single master roles, since this is a single domain forest.

\*\* These two applications are now hosted on two separate Windows 2000 servers.

\*\*\* One of the file and print servers is a Windows 2000 computer with the WINS service.

\*\*\*\* Each regional office has a single Windows 2000 server that serves both roles.

## Analyze Security Considerations

The security requirements defined in the business analysis should be compared with the current security configuration. If the security requirements dictate changes in the present security configuration, you will need to determine if the network resources can handle

the load of whatever security practices you are putting in place. For instance, if the authentication method will add a certain network load, you need to determine if you need to be using, say, IPSec. If this is the case, you need to be aware of the additional traffic it puts on your network.

The tools mentioned earlier, LoadRunner and DynamicAccess, are useful in determining the load that security places on the network. Other tools also help determine if you are achieving the level of security required in the business analysis. The following list is a sampling of security analysis tools:

- SAFeSuite Decisions by Internet Security Systems ([www.iss.net](http://www.iss.net)) combines their individual security assessment products and includes Internet Scanner, System Scanner, Database Scanner, and RealSecure. This suite allows you to use these applications and select third-party firewall and intrusion detection systems as a combined analysis tool.
- SecurityAnalyst 5.0 by Intrusion.com ([www.intrusion.com](http://www.intrusion.com)) is an assessment tool that provides centralized audit data of all key Windows 95/98/NT/2000 and Novell NetWare security features. It analyzes six critical security areas: password strength, access control, user account restrictions, system monitoring, data integrity, and confidentiality.
- BindView's bv-Control ([www.bindview.com](http://www.bindview.com)) is an enterprise-wide security assessment tool for Windows 2000 and NT. From a single console, bv-Control for Windows 2000/NT continually audits configuration standards and operational performance across multiple domains and alerts the system administrator when risks are identified.



We have provided URLs for the vendors of the products listed here and elsewhere in this book. You should explore these sites, because most of these vendors have additional, related products that are not listed here, and they all update their products to add functionality and compatibility with new operating systems.

---

## IMPACT OF DESIGN ON EXISTING AND PLANNED TECHNICAL ENVIRONMENT

No man is an island, no dog howls alone, and no network change exists in isolation. Every nuance of an infrastructure design—and every change to that design—affects something else. A smart network designer figures out the impact in advance.

Microsoft's objectives for Exam 70-221 list the following as part of the analysis of the impact of infrastructure design on the existing and planned technical environment. We think it is a pretty comprehensive list.

- Analyze network infrastructure, protocols, and hosts.
- Detect current applications and their impact.

- Analyze network services.
- Analyze TCP/IP infrastructure.
- Analyze current hardware and performance.
- Identify existing and planned upgrades and rollouts.
- Analyze technical support structure.
- Analyze existing and planned network and systems management.

We discuss each in turn in the following sections.



When planning a network infrastructure design, you will use a discovery process in which you should not only examine the additional network requirements but also look for inefficiencies in the current network usage. This is an opportunity to fine-tune the existing network so that you get the best performance from it before you add additional load and resources. (This alone could make you a hero.) You must always be on the lookout for applications that are running unofficially on the network, and which are not required or desired by the organization. You must also watch for both unnecessary and inefficient use of protocols.

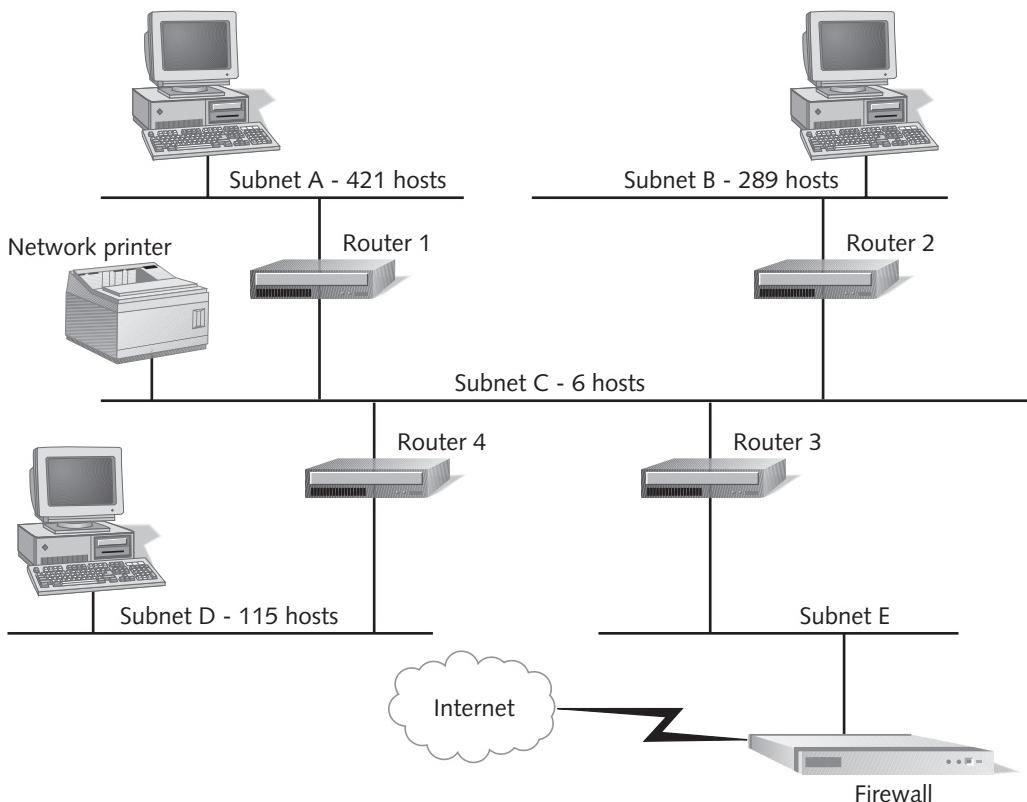
## Analyze Network Infrastructure, Protocols, and Hosts

When analyzing the network, you will be examining the existing infrastructure protocols and hosts. You should start with an inventory of the physical infrastructure and the protocols in use on the network on both server and client computers, looking for inefficiencies that are wasting bandwidth.

### Network Infrastructure

Inventory the Physical and Data Link layer network infrastructure, including all existing routers, bridges, switches (indicate the type), and so on, as well as the media. If you are dealing with a small network, you can probably do an actual hands-on physical inventory, but if the network is large or geographically dispersed, it is impractical to go to each location to physically look at each piece of equipment.

At this stage, it is reasonable to ask for a diagram, which should also include the hardware make and model of each computer and its location. The difficulty of acquiring this information is inversely related to the budget and/or the technical sophistication of the IT department. Some IT managers will have current and accurate information at hand; others will struggle to gather the information. Included in this diagram should be all WAN links and the type of service provided over each WAN link. This information is important because it is the foundation for your entire network design. A network diagram might be as simple as Figure 3-1, which showed a high-level overview, or more complex as in Figure 3-9, which shows the headquarters network for Novl Gifts.



**Figure 3-9** Headquarters network for NovlGifts

Next, perform an analysis of the present usage, looking for inefficiencies that can be eliminated, such as inappropriate usage by users downloading large files from the Internet that are not for work-related tasks. See Figure 3-10 for more examples of network inefficiencies.

The following discussions of protocols and applications include tips on further optimizing the use of the present network infrastructure. After unnecessary traffic has been eliminated from the network, you can determine how the network infrastructure must be upgraded for the new protocols and services.

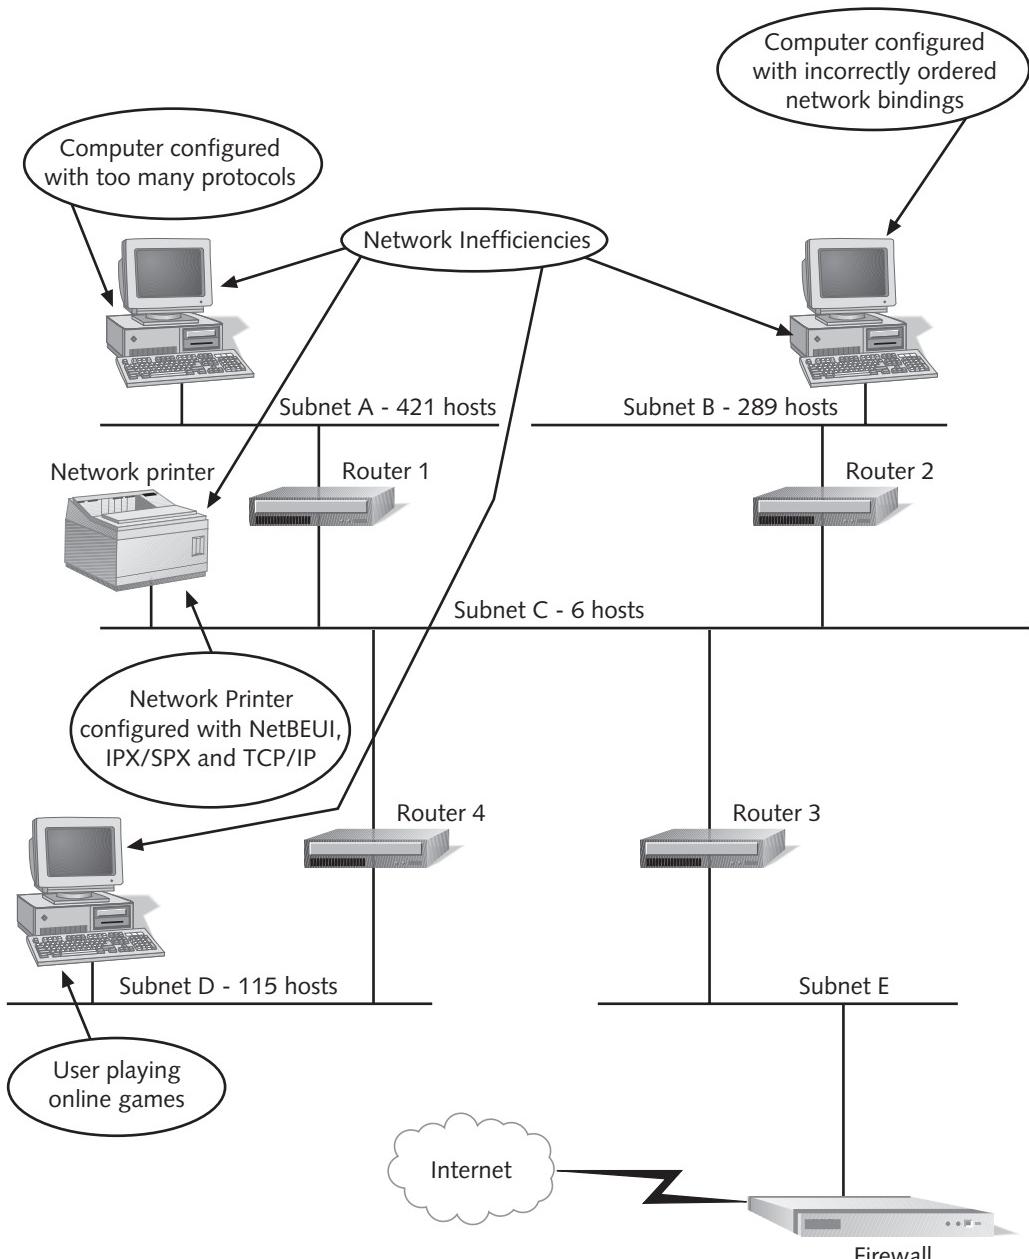


Figure 3-10 Network inefficiencies

### Protocols

*The fewer the better!* This is the first rule of protocols and it should be applied to protocol suites such as TCP/IP, NetBEUI, Apple Talk, IPX/SPX, and DLS (a protocol used for communicating with IBM mini- and mainframe-computer networks). One way to

avoid having to add new network infrastructure is to eliminate inefficient use of the existing infrastructure. Always be on the lookout for unnecessary protocol suites. We call these rogue protocols, although you might say that the person who allowed these protocols to exist on the network might be the real rogue.

One example of the creation of rogue protocols occurs with the use of some network print devices. Some of these devices will, by default, enable many protocol suites. Fortunately, the use of a Network Monitor capture will show you the network traffic generated just by the printers talking to each other in the protocols!

Spanning Tree Protocol also is a rogue protocol. It is enabled by default on many network devices, but it is only necessary when switches are connected together with redundant paths. If your switches aren't configured this way, you're just broadcasting bridge protocol data units (BPDUs) needlessly.

Another place to look for unnecessary protocols is on servers and client computers. Since Windows products can have many protocol families installed simultaneously, you need to eliminate any protocols not actually needed by the servers and clients!

The following actions help find unnecessary protocols:

- Using a network packet analyzer, such as Sniffer from Network Associates or Microsoft Network Monitor, you can capture network traffic at various times and identify all protocols, their sources, and their destinations. Then, you can investigate why they are being used.
- While examining network traffic, you can look for the percentage of broadcast traffic. Consider 20% broadcast traffic to be the upper limit.
- On every computer (server and client), examine the network configuration and identify all unnecessary protocols and remove them.

There are network protocol configuration missteps that can cause unnecessary traffic. Microsoft NetBIOS clients have a setting called “node type” which controls how that computer goes about resolving NetBIOS names to IP addresses. The node types are B-Node, M-Node, P-Node, and H-Node. For detailed information on node types, go to Microsoft’s Technet site ([www.microsoft.com/TechNet/](http://www.microsoft.com/TechNet/)) and search the Knowledge Base for article Q119493.

We will not go into too much detail on node type here, except to remind you that a B-Node NetBIOS client broadcasts to resolve NetBIOS names. This not only generates unnecessary traffic, but also can result in failure to resolve names in a subnetted network. So, all the extra traffic of broadcasts would be for naught! The desired node type in a subnetted network with WINS servers available is H-Node (Hybrid Node). An H-Node client first checks its cache of recently resolved NetBIOS names. If the name is not found in cache, the client then queries a WINS server. If the WINS server cannot resolve the name, the client performs a NetBIOS broadcast. A statically configured Windows 2000 WINS client will be H-Node by default. A DHCP client must have node type configured in scope options on the DHCP server.



Using Microsoft's Network Monitor, a consultant we know once found 62% broadcast traffic on the network of a large organization. Most of this was caused by network printers and computers configured with too many protocols. Just by doing the simple tasks listed here, he reduced the broadcast traffic to 8%!

In some cases, multiple protocols are indeed needed to communicate with various network resources. In these cases, apply the second rule of network protocols: *Make sure the bindings are in the correct order!* This requires that you verify that the network bindings on each client and server are in the order in which they are required, from most frequently used to least frequently used.

Having the most frequently used protocol first in the binding list will reduce the average connection time. For example, if NWLink and TCP/IP are both installed on a computer, and if most servers that the computer connects to are using TCP/IP, you will want to make sure TCP/IP is first in the binding order. Figure 3-11 shows the Adapters and Bindings tab of the Advanced Settings dialog box. In this case, a network interface has both TCP/IP and NWLink bound to it. If the TCP/IP protocol is used more than the other, you would make sure it is first in binding order, using the arrows on the right to change the binding order. If NWLink is installed on this computer, but not used by this interface, then you would clear the check box to disable that protocol on this interface.

Remember that some protocols are faster than others for certain network topologies. If two protocols are equally used, but one proves faster than the other, then move the faster protocol to the top of the binding order to improve performance. You can try Hands-on Project 3-5 to learn more about protocol binding order.

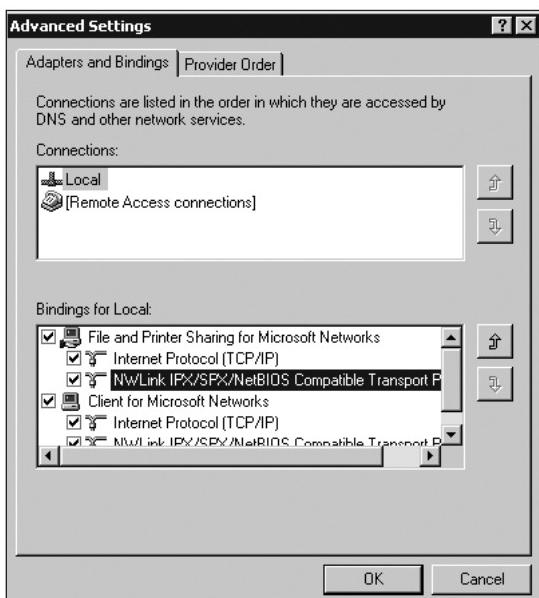


Figure 3-11 Network protocol bindings

## Detect Current Applications and Their Impact

A current application is one that is presently in use on the network, but not necessarily “new”—it could even be a 12-year-old DOS application that someone loves and still uses. It is not unusual for an organization to use an old application because nothing else is available for that particular function.

Several years ago the authors were working with a healthcare organization that identified more than 200 DOS applications in their various nursing homes, clinics, and care centers. The majority of these were “vertical market” applications, which are applications written for the needs of a certain industry. Many were homegrown or created and sold to them by companies that were no longer in business. The search for consolidation of what they were using and replacement applications took years of work.

Do not be surprised when you find such applications, and do not be too hasty in disposing of them until you verify that there is a satisfactory replacement. All of this may not seem truly part of your job. However, when you ask the right questions, you will discover things that others have simply accepted and lived with.

You want information about these existing applications because the applications in use on the network all contribute to the load on the network. If you can initiate a movement to remove unnecessary or outdated applications from the network, you will discover “found bandwidth”—bandwidth that was misused.

Information about what applications are currently in use in the organization can be gathered in interviews, from physical audits or from automated audits. This is part of the multistage process of building the application portfolio. After the initial interviews, you will find yourself returning to the managers as you discover more applications through the audits. This allows you and the managers to refine the list of applications and services needed.

Software audits are good, but just knowing you’re running Exchange (for example) doesn’t tell you much. For instance, one network might average 2000 e-mails per day with 10 MB attachments, while another might average 20,000 e-mails per day with no attachments. The server capacity tools are probably sufficient for server resources, but bandwidth planning should be done from statistics taken directly off the wire with a hardware-based protocol analyzer or **RMON** probe, or at the very least, from the utilization statistics generated by your network switches and routers. Although the switch and router statistics aren’t as accurate as the protocol analyzer and RMON probes, they are still more accurate than PC/software-based tools like NetMon because they operate at the speed of the hardware and aren’t subject to the whims of the Windows operating system. Even so, these tools are expensive, and bandwidth planning with NetMon is better than no bandwidth planning at all.



Remote Monitoring, or RMON, is a network management protocol that allows network information to be gathered at a single workstation from network hubs and switches that support the protocol.

You need software audits because the software inventory information gathered in interviews is partially anecdotal, and it needs confirmation before you use this information in your design. You may be lucky enough to work for an organization that stays on top of their software inventory, but just in case, we suggest tools such as Microsoft's Systems Management Server (SMS) or Blue Ocean Software's Track-It!

If your audit reveals applications that are not on the list compiled from the business analysis, you need to have management determine if they are necessary to the business. If they are necessary to the business, but were simply overlooked in the business analysis, then add them to the applications portfolio. However, if you find applications that are not deemed necessary to the business, then someone must make a decision concerning the future use of these applications on the network. If they are allowed to persist, then you will still need to add them to the applications portfolio and include them in the Load bucket of the Iron Triangle.

There are several tools available for performing software audits. Software audits can help you, as a network designer, to determine what really is running on the network. These tools can also help you ensure that you are running legally licensed software, which, although not directly tied to network design, is very important for any organization. A large organization with the necessary budget and resources might be using Microsoft SMS for software inventory. If this is the case, you may be in luck. If a recent software inventory has been successfully taken, you may simply need to request that the SMS administrator run a report for you.



SMS is a Microsoft Back Office product with many facets, including hardware inventory of desktops and servers running Microsoft operating systems, software inventory of those same systems, software metering for the enforcement of site licenses, an enhanced Network Monitoring tool, software distribution, and remote tools for remote support of Windows desktops. It is a *huge* product, one that is usually implemented in large organizations that can provide the dedicated staff or expensive consultants to install and administer it. From a support standpoint, it is bigger than even Microsoft's mail server product, Exchange. It's just not as visible.

There are third-party products that provide software inventory capability. One such product is Track-It! from Blue Ocean Software ([www.blueocean.com](http://www.blueocean.com)). They offer a demo from their Web site.

The Software and Information Industry Association (SIIA), formerly the Software Publishers Association ([www.spa.org](http://www.spa.org)), has free tools that will help in ferreting out just what applications are out there on your client computers. From their home page, click

the Anti-Piracy link. On the Anti-Piracy page, click the Asset Management Tools link. At the time of this writing, they had two tools available at their Web site:

- WRQ Express Inventory Version 4.5, SPA Edition from Express Metrix ([www.expressmetrix.com](http://www.expressmetrix.com))—A small organization with a limited budget might consider WRQ Express, which is a free limited-license version; you may purchase the full version from Express Metrix. The SPA Edition of WRQ Express Inventory is a software tool that performs a one-time audit on up to 100 personal computers within a 30-day period. The product includes a summary report titled, “Am I Legal?”
- KeyAudit 5.0 from Sassafras Software ([www.sassafras.com](http://www.sassafras.com))—KeyAudit is distributed by SIIA as part of its software anti-piracy educational efforts. KeyAudit searches for all Windows or Macintosh applications and will produce indexed lists. Sassafras also offers the commercial product, KeyServer, which enforces software licenses.

## Load Created by Current Applications

Once you have eliminated unnecessary protocols and applications on your network and have determined the contents of the current application portfolio, you will need to analyze the load that the current applications place on your servers and network.

For measuring server load and determining server capacity, there are these tools:

- BMC Software BEST/1 ([www.boole.com](http://www.boole.com))
- SAS Institute’s IT Service Vision ([www.sas.com](http://www.sas.com))
- TeamQuest’s Model ([www.teamquest.com](http://www.teamquest.com))
- Metron’s Athene ([www.metron.co.uk](http://www.metron.co.uk))
- HyPerformix ([www.hyperformix.com](http://www.hyperformix.com))

For measuring network load and determining network capacity, look for tools from:

- Compuware ([www.compuware.com](http://www.compuware.com))
- CACI ([www.caci.com](http://www.caci.com))
- MIL3’s OPNET Technologies ([www.mil3.com](http://www.mil3.com))
- Optimal ([www.optimal.net](http://www.optimal.net))
- Make Systems’ NetMaker MainStation ([www.makesystems.com](http://www.makesystems.com))

Stress testing is a time-honored technique for capacity planning. When you stress test, you simulate the load that your planned applications will have on your server and network. You should definitely plan to use stress-testing tools such as:

- Mercury Interactive’s LoadRunner ([www-svca.mercuryinteractive.com](http://www-svca.mercuryinteractive.com))
- RSW Software’s e-Test Suite ([www.rswsoftware.com](http://www.rswsoftware.com))

- Segue Software's Silk product family and eConfidence programs ([www.segue.com](http://www.segue.com))
- RadView Software's WebLoad ([www.radview.com](http://www.radview.com)) (also licensed by Computer Associates)
- Loadtesting.com's Portent and Portent Supreme ([www.loadtesting.com](http://www.loadtesting.com))
- Rational Software's Rational Suite TestStudio and other products ([www.rational.com](http://www.rational.com))

## Capacity Assessment

Good network design implies knowledge of current capacity. No matter what tools you use for capacity planning for your network, network connection devices, and servers, you must do the following:

- Identify trends in usage, set up logging to files, and gather historical data at various intervals.
- Monitor real-time performance data. For ongoing management of the network, set administrative alerts based on thresholds.
- For historical data, analyze the logged performance data, looking for peak usage periods.
- For capacity planning, use the performance data collected previously to predict performance under various scenarios. This can be accomplished with some of the tools listed in this chapter.
- Tune the performance. This includes elimination of unnecessary protocols and applications on the network, as well as an understanding of how each service and protocol can be configured for optimal performance.

## Analyze Network Services

Your next step is to determine which network services are currently in use and how they are configured. You can use one of the network assessment tools mentioned in this chapter to analyze the load that network services are currently placing on the network. Network services include:

- IP address allocation (DHCP)
- Authentication
- Active Directory services
- File and print services
- Web services
- FTP services

- Name resolution services
- Database services
- Network management services
- Remote Access services
- Routing services

In your analysis, determine if any of the present services will be retired. For example, if all the NetWare servers are being removed from the network, then any routing services to support IPX/SPX can be removed.

Next, determine whether the remaining services are being utilized as efficiently as possible. For example, your current WINS servers may be Windows NT servers. However, the Windows 2000 implementation of WINS server has been greatly improved, both to be easier to manage and to perform faster replication with WINS replication partners. So even though you are migrating to a Windows 2000 Active Directory, if you continue to have legacy operating systems and legacy applications on the network, you will continue to need WINS for NetBIOS name resolution. For this reason, consider upgrading WINS servers to Windows 2000.

In looking at network services, consider how the new design will affect the functionality of any of the current network services. If this is part of a migration from Windows NT to Windows 2000, the downlevel clients will still use their old authentication methods. That is, Windows 9x clients still will use LAN Manager authentication, unless they have the Directory Service Client installed (dsclient.exe). In this case, they may use NT LAN Manager V2. Windows NT clients will continue to use NT LAN Manager or NT LAN Manager V2.

In contrast, only Windows 2000 clients will use the new Kerberos authentication method for user and computer authentication. This means that you will still have the old authentication methods on the network, as well as a new authentication method. The network traffic patterns will vary by protocol. Once again, Network Monitor (or third-party tools) will help in analyzing the load that the various authentication methods place on the network.

It is important to compare your list to the manager interviews you did for the business analysis so that you can establish the relative importance of each of the services. Some, such as DHCP, directory services, and name resolution services, are critical to basic network function. Others, such as file and print services, Web services, FTP services, database services, and network management services, might provide you with some “wiggle room” if the network design becomes too expensive to implement in its entirety, or if implementation must be spread over several months to spread out the cost.

## Analyze TCP/IP Infrastructure

During your technical analysis, gather information about the current TCP/IP infrastructure to compile an overall inventory of the protocols and services in use and the

strategy presently in place for utilizing IP addresses. This inventory will become your composite view of the existing TCP/IP network. You may build upon this composite view or you may have to rip it apart to accommodate the new design. These are the topics you need to find in this part of the analysis:

- Addressing strategy and address assignment
- Naming strategy

3

## Addressing Strategy and Address Assignment

The strategy used to allocate addresses throughout the enterprise network can have major implications, especially if there was a poorly planned strategy or no strategy at all! Without careful planning and establishment of a strategy, you risk the following problems:

- Duplicate addresses
- Illegal addresses
- Too few addresses, or too few for subnets with a larger number of hosts
- Wasted addresses

You should analyze the current addressing strategy, looking for use of the following recommended practices:

- A plan for meaningful addresses includes the practice of assigning certain ranges of addresses within each subnet to classes of devices. For instance, the policy might state that on each subnet, all host addresses from 1 to 10 are used strictly for routers, all addresses from 225 to 254 are used strictly for servers, all host addresses from 200 to 224 are used strictly for network printers, and all addresses between are used for client hosts. Be sure to analyze the maximum number of each category of device needed on any subnet, plus allow for growth in the number of such devices. A variation on this strategy is one in which the IP addresses of network devices fall on binary boundaries.

Meaningful address policies are enforced so strongly within some organizations that the authors have encountered network administrators who were convinced that it was mandated by an RFC. It is just a sign that someone did some effective planning for how IP addresses were to be used.



These practices will seem brilliant to you if you use network analyzer traces, because you will recognize the type of device generating the traffic by its name, provided your network analyzer resolves names, which the SMS version of Microsoft Network Monitor can do.

- Hierarchical IP addresses involve assigning one IP network address, such as 10.0.0.0/8, to your organization's intranet and subnetting it in a hierarchical manner for IP addressing. This will facilitate network management, because network maps will be logically, as well as physically, organized. With the use of

variable length subnet masks (VLSMs), this strategy also makes it easier to allocate logical subnets containing the correct number of hosts for both present and future usage. For instance, if you only need two host addresses on a subnet, as in the case of a point-to-point connection, you could subdivide the network address 10.0.0.0/8 with a network address of 10.0.0.0/252, which would give you a subnet with two hosts. This further division of a network address with subnet masks of different values is the essence of VLSMs. And, in case you still aren't sold on this method, this strategy also makes it easier to use network filters on firewalls, routers, bridges, and switches to eliminate unauthorized traffic for network optimization and security.

- It allows room for growth. A good plan allows for growth in the number of addresses needed, both per subnet and enterprise-wide. To do this well, you analyze the entire IP structure, looking at the number of existing and planned addresses needed per subnet. Then, using the hierarchical strategy, and using variable length subnet masks (VLSMs), you can assign logical subnets to each subnet that will allow for growth, with minimum wasted addresses.
- It allows dynamic addressing wherever practical. Dynamic Host Configuration Protocol (DHCP) is one of the best laborsaving devices in the TCP/IP suite of services and protocols. It is a no-brainer to use DHCP to give out addresses to client computers. We normally assign addresses to servers manually; however, do consider using the reservation feature of DHCP to reserve IP addresses for individual servers, based on each MAC address.
- It allows private addresses for security when appropriate. There are several ranges of IP addresses that are officially specified in an RFC by the IETF as private addresses. This means that they are to be used exclusively on internal private networks. In fact, if an Internet router detects a packet with one of these addresses, the packet will be discarded. The addresses reserved for private addresses are:
  - 10.0.0.0/8—A class A private network address that includes a range of addresses from 10.0.0.1 to 10.255.255.254. It has 24 host bits, and provides the greatest number of subnet and host configurations.
  - 172.16.0.0/12—Is actually a range of 16 class B networks. It consists of all the addresses in the range of 172.16.0.1 to 172.31.255.254.
  - 192.168.0.0/16—Is a range of 256 class C network IDs, consisting of all the addresses in the range of 192.168.0.1 to 192.168.255.254.

Hosts using these addresses on a private network that want to access resources on the Internet must go through a device that provides Network Address Translation (NAT) or use an application layer gateway, such as a proxy server. In both cases, the host providing the address translation or gateway must have an interface with a valid Internet address.

## Naming Strategy

You will need to have structured host naming, name registration, and name resolution strategies. In a Windows 2000 domain, using DNS is the rule for the name space. Your overall domain naming strategy is part of your Active Directory design, which is beyond the scope of this book. If the Active Directory domain structure is already in place, this is certainly a “done deal.” However, if it is still in the planning stages, find out what the strategy for domain naming is and whether the Active Directory planner included a naming strategy for the hosts and shared printers on the network.

*Try to work with the established strategy for naming.* To simplify things, we will refer to the strategy of naming these devices as the “host-naming strategy.” However, our host-naming strategy does not involve only DNS. We cannot overlook the NetBIOS namespace in our plan. If you continue to have legacy operating systems (Windows 3.x, Windows 9.x, and Windows NT) and legacy applications on the network, you will continue to need WINS for NetBIOS name resolution.

*Always consider the lowest-common-denominator rule of name resolution.* If the client operating system is a legacy system, the name space the client defaults to is NetBIOS, even in a Windows 2000 domain. If the client operating system is Windows 2000 authenticating to an NT domain, then the default name space is NetBIOS. When a Windows 2000 client authenticates to a Windows 2000 domain, it uses DNS to locate the domain controller for authentication.

Unless you have purged your network of legacy clients and legacy applications, you will need to maintain the two name spaces. Therefore, your plan must accommodate both name spaces. Keep machine names under the 15-character limit of NetBIOS, in spite of the fact that Windows 2000 can identify a separate DNS and NetBIOS name.

The best host-naming strategies use short names that have meaning to the users, uniquely identifying the device. The names might follow a standard that can identify the device, such as using “srv” in the name for a server. Table 3-3 lists common naming codes.

**Table 3-3** Resource naming codes

Abbreviation	Type of Device
Srv	File and print server
Web	Web server
Cli	Client desktop computer
Cln	Client notebook computer
Ptr	Printer

Names such as those found in Table 3-3 would suffice if a department, such as accounting, is known to be in one location, but if there are several locations where accounting

servers are located, have a policy that includes location as well. To implement this policy, you might have two- or three-letter abbreviations for each office location.

This naming specificity can be carried to whatever level is needed. For example, you can have three characters to indicate the department, based on official corporate abbreviations. The next three characters could indicate the business unit, also based on official corporate abbreviations. The final three characters could indicate the number of this type of device within the department. Thus, “bossrvactfro002” would be the second Boston-located file and print server in the accounting department of the Frozen Foods business unit.

Some other naming tips include the following. These tips will greatly reduce troubleshooting times, especially with novice or new administrators.

- Make the physical label of printers include the server and share name of the printer so it's easy to find.
- Make your printer name include the model of the printer.
- In your closets, label the end of the patch cable that's plugged into the switch with the name of the jack on the wall, and label the end plugged into the jack with the slot and port number of the switch.
- Use Visio to print pictures of your switches, then put them in plastic protector sheets and write user and server names next to the ports in grease pencil.
- Use names like bossrvactfro002 INSIDE and use nonsensical names like Ren and Stimpy or Yakko, Wacko, and Dot for OUTSIDE or registered Internet services.

Whether an organization has either a centralized or decentralized technical administration, the design of the strategy should be accomplished centrally and include details on how ongoing administration will be accomplished.

## Analyze Current Hardware and Performance

At this stage, you need to inventory current hardware and usage. You can't determine what you need to acquire in the future if you don't know what you currently possess. For client and server hardware inventory, use Microsoft's SMS, if it is used and in place, or third-party software such as Track-IT!, mentioned earlier in this chapter.

Gather performance data on the existing NT and Windows 2000 systems by using logs created by Performance Monitor or one of the third-party tools listed in the Software Load Testing section earlier in this chapter. Determine if these servers can or will be used for any additional services identified in the design. If so, test whether they can handle the additional load. Once again, consider one of the load-testing products listed earlier in this chapter. After the tests are completed, define what upgrades must be done to existing equipment and what additional equipment must be purchased.

## Identify Existing and Planned Upgrades and Rollouts

In the business analysis, you compiled the list of applications that should be included in the application portfolio. In your technical analysis, you identify which of these applications is scheduled for an upgrade. In addition, you identify which of these applications is entirely new to all or part of the user community and will involve a rollout.

You will need to determine the schedule for the upgrades and rollouts, and verify that each will be tested on a test network before being introduced to the production environment. In addition, you need to verify that the test environment will include any aspect of the new network design (services, protocols) that could affect the functionality of the applications.

As scary as it may be to consider application upgrades and application and operating system rollouts on top of a network redesign, it may not actually be as bad as it sounds. This may enable you to eliminate excess protocols on the network that are used by older applications and operating systems. Windows 2000 clients and newer applications will take advantage of more features of the Windows 2000 network environment you are designing.

## Analyze Technical Support Structure

Much of your technical support structure information was gathered during the business analysis of Chapter 2. You will use that information to determine if technical support can perform the upgrade under their present structure, and if they can support the new network infrastructure after it is in place. The cost of bringing in additional people for implementation and support must be included in the plan.

A technical support structure is not monolithic. It has these easily identifiable components:

- The number of people in the support group, their education, their years of experience, their specialties, their certifications, the special training they have had, and the shift that they work
- Whether anyone is planning on leaving
- Whether anyone has training experience
- Whether anyone has lived through a previous design change

It is important to understand the capabilities of the present support people and match them against the needs of the new network. If you are introducing protocols or services with which no one in the support group has experience, you will have to either train someone or bring in outside help. If you have employees with needed experience and they also happen to have some training experience, you need to determine whether you can use them to train other support people. You also need to analyze whether the total number of people you have and their geographic distribution is sufficient to install and support the new network. If not, you will have to expand the group. All of this costs money, which must be included in the overall costs of the new network design.

## Analyze Existing and Planned Network and Systems Management

At this stage, you need to determine if the organization has adequate procedures and tools in place to manage the existing network and systems. If it does not, you will need to add procedures and tools to support the new design.

If changes need to be made to existing procedures, assign someone the task of documenting these procedures and include in your plan the training of personnel to learn these procedures. For instance, if you will be using the free 3Com tools to measure the network performance through your 3Com connection devices, determine that administrative tasks—such as setting up logging of activity and alerts, establishing archives of performance information, generating reports, and reviewing the data to search for trends—are assigned appropriately.

If changes need to be made to network and server management tools, start shopping. Several of the vendors mentioned earlier in this chapter have network and systems management tools. Microsoft's SMS is a Windows management tool with software and hardware inventory capabilities, remote tools (including remote control), software distribution, and software license metering. IBM, Compaq, and Hewlett-Packard have tools for network management. Many vendors, including Microsoft, offer free demo software for evaluation, downloadable from their Web site and/or available on CD. (Blue Ocean Software delivered a Track-It! demo CD to the authors of this book within four business days wrapped around a holiday weekend!)

---

## ANALYZE NETWORK REQUIREMENTS FOR CLIENT COMPUTER ACCESS

Sure, you know that you need a new network design to meet the high-level goals of the organization. Nonetheless, you still have to take the current needs of the end users into account because, like it or not, the entire purpose of the network is to serve the end user. Regardless of who your end user is, you cannot thwart his or her work goals.

At this stage, you need to look at how the user accesses the network. The analysis of the network requirement for client computer access (end users, for the most part, use client computers) has two parts: the analysis of end-user work needs and the analysis of end-user usage patterns.

### Analyze End-User Work Needs

For this part of the analysis, look at what network resources each user must access and where these resources are located in relation to the user. For instance, you will find users who access file and print servers in their own LANs, but who access an e-mail server across a WAN link. Other users may be accessing all resources across dial-up connections. All of these permutations must be documented.

Just as a network account administrator works to group users into security groups or distribution groups for ease of administration, as a network planner, you will analyze and plan for end-user work needs by grouping users into groups with common network access needs. This may involve a reworking of some of your groups, but will pay off in giving you a clear analysis of the work needs. Table 3-4 shows how Novl Gifts employees (in security groups) access resources across the WAN.

**Table 3-4** End-user work needs by security group

Security Group(s)	Number of Users	Resource Used	Resource and Location
Acct_receive	20	Accounting database Home directories	Unixsrv1 – headquarters F&PSrv1 – headquarters
Acct_payable	15	Accounting database	Unixsrv1 – headquarters
Acct_payroll	13	Accounting database	Unixsrv1 – headquarters
Order Entry	35	Inventory database	Unixsrv2 – headquarters
Sales & Marketing staff	55	Inventory database	Unixsrv2 – headquarters

## Analyze End-User Usage Patterns

End-user usage patterns include patterns of network resource access. Look for daily, weekly, monthly, and even yearly patterns of network usage. Recall the business analysis in which you determined whether the company had any product or service cycles. The data you gathered will lead you to the usage patterns. For each network resource, look for the following:

- When is the resource accessed?
- What is the average length of a connection to the resource?
- How much network traffic is generated per connection session?
- How many end users are simultaneously connected to the resource?
- What are the security requirements?
- How much bandwidth is needed for security?
- What is the calculated bandwidth needed for this usage?

---

## ANALYZE DISASTER RECOVERY STRATEGIES FOR THE EXISTING TECHNICAL ENVIRONMENT

Your analysis of the existing disaster recovery strategy of the technical environment is only a piece of the comprehensive disaster recovery strategy that every organization

should have. A technical disaster recovery plan for a Windows 2000 computing environment will involve three major areas:

- Disaster recovery strategies for client computers
- Disaster recovery strategies for servers
- Disaster recovery strategies for the network

The disaster recovery strategy for client computers, servers, and the network should include an audit of current procedures. In addition, there should be written procedures for backup of all data and recovery of all systems in the event of a disaster.

All procedures should be tested on a regular basis. A theoretical plan that has not been tested is fraught with potential for failure. The authors of the plan will find that instructions are often open to interpretation, and steps may have been omitted. The disaster recovery plan may even be based on flawed assumptions, like the capacity of the tape backup systems, and the plans could be out of date and not even cover new servers. Fortunately, these flaws and missteps can be discovered through a test recovery, and corrective measures can be taken. These measures could include simply rewording the written instructions or may require revamping some or all of the plan.

A disaster recovery plan should include provisions for using manual procedures for any processes that can be performed manually. Your aim is to be back up and running as soon as possible. Thus, you can still have orders taken manually by a sales force that visits client sites, for example. Their manual efforts will buy you time as you fix the network.

If a comprehensive disaster recovery strategy is not defined in a formal policy, it is time to start asking more questions. For starters, try these:

- How much money would the business lose if the network were down for one day, two days, etc.?
- What business processes can be achieved without the network, and do people know how to do the manual processes?

The most stringent disaster recovery plan should be defined for the network components that support the most critical applications and services in your organization. Look for the need for fault tolerance, such as hardware-level RAID on servers, server clustering, and redundant WAN links.

## Disaster Recovery Strategies for Client Computers

Client computers are those computers to which users have physical access and at which they accomplish their work. Some client computers have standard office productivity tools installed directly on them. Others may use software that is run from the server. No matter the client computer configuration, if you are concerned about recovering users' data, the client computer should be considered an access tool, not a place to store data. That way, in case of disaster, the client computers can simply be replaced by computers

with the appropriate configuration for the users, without concern for locally stored data (because company policy states all data must be stored on servers, and client computers will be reimaged after a failure).

There are now creative and flexible options for recovering the client operating system, applications, and configuration. One option is to have a server-based source from which automated installs can be run over the network. Another option is to use a third-party imaging tool and keep client images on a server from which images can be brought down and installed on client computers.

You also can have client computers that are actually just dumb terminals connecting to mini or mainframe host computers or that are running terminal emulation software to connect to these larger systems. Windows 2000 includes Terminal Services, which provides multi-user access to a Windows 2000 server, in which several users can run sessions simultaneously on the server from their computers. All application processing occurs at the server hosting the terminal services.

Disaster recovery for terminal services clients focuses on the servers. Recovering the servers is like recovering many user desktops. Although you will still have to have a plan for recovering the client computers, if the users are only accessing Terminal Services from their client computers, recovery of the clients should be much less involved than a more conventional client. You simply restore the operating system and install the Terminal Services client on the client computers.

Prevention is also part of disaster recovery. Good security policies can prevent disasters. Do not allow unauthorized physical access to user computers if it can be avoided. Use the highest security authentication protocol available for the client computer. If the client is a Windows 2000 computer logging on to a Windows 2000 Active Directory, it will by default use Kerberos authentication, which is far better than the authentication protocols of Windows NT 4.0 or legacy Microsoft clients. Even the best authentication protocol is useless if you do not require good security practices.

Recommended practices for disaster recovery of client computers include the following:

- Have a well-publicized policy of no data stored locally.
- Have a well-publicized policy of doing complete reimages or system replacement if there is a system failure—individually or as part of a larger disaster.
- Have a well-publicized policy for physical access to user computers.
- Have a well-publicized policy for strong passwords.

## Disaster Recovery Strategies for Servers

The fundamental disaster recovery tool for servers is a backup strategy that is adhered to and tested on a regular basis. In addition, hardware-level RAID 5 and disk mirroring should be employed for fault tolerance. With critical systems, carry the fault tolerance further with server clustering, which will be discussed in later chapters. The icing on the cake

would be a completely redundant data center located at a distance from the existing data center, but ready to go online on short notice to provide the most critical applications in case of a large disaster. Check to see if your organization has this strategy in place.

## Disaster Recovery Strategies for the Network

Of course, your efforts to restore clients and servers are for naught if you don't restore their environment as well. You must, in advance, identify existing disaster recovery strategies for the network, such as the following:

- Redundant WAN links
- Guarantees from your network WAN provider(s) that the links are truly redundant (if possible, have separate providers for these redundant links)
- Multiple routes between users and network resources
- Elimination of single points of failure
- A formal, written disaster recovery plan, including scheduled tests

Don't forget your power source. Use redundancy and/or make a major investment in a large, uninterruptible power system. One airline lost many hours of reservations system access, and associated business, not too long ago because the redundant power source to their network passed through the same physical conduit as the primary source. During excavation for construction of a new runway, both lines were severed.

---

## CHAPTER SUMMARY

- In this chapter, we evaluated technical requirements. The first section of the chapter gave some insights and tools to help you analyze an organization's existing and planned technical environment and goals. This time the tools were not only questions to ask, but also actual software tools that can be used for a variety of data-gathering efforts. We provided references to a number of tools you can use to actually measure the existing network usage and loads and some insights into what to look for.
- The second section of this chapter gave you more insight into the huge impact that infrastructure design has on both the existing and the planned technical environment. We gave you ways to measure the real network load and ways to identify whether the load is "real" or whether it is made up of unnecessary overhead network usage.
- The third section focused on helping you discover the network requirements for client computer access. You not only need to discover end users' work needs, but also their work patterns, because that heavily affects the network usage.
- Finally, we focused on learning about the existing disaster recovery strategy for client computers, servers, and the network. Many organizations have only primitive disaster recovery schemes or don't practice disaster recovery at all. But disaster recovery is a critical part of network infrastructure design. If it doesn't exist, the organization is in real trouble.

---

## KEY TERMS

**bandwidth**—The amount of data that can be transmitted in a fixed amount of time, usually expressed in Kbps or Mbps.

**bindings**—Define the relationships between networking software components. By default TCP/IP, NetBEUI, and NWLink, if installed, are bound to all network interface drivers.

**cloud**—Jargon used to describe a network where a given packet could take one of several paths to get to the destination. It's the lack of visibility (the inability to know which path will be taken).

**domain-naming master**—A forest-wide single master operations role that is automatically assigned to the first domain controller in the first domain in the forest. The domain-naming master manages the addition and subtraction of domains in the forest.

**floating single master operations (FSMOs)**—The original term used for single master operations. This term is still used in documentation and in tools that let you move the roles, such as the utility NTDSUTIL.

**global catalog server**—A special role for one or more domain controllers in a Windows 2000 Active Directory domain. The global catalog server contains a partial replica of every domain directory partition in the forest as well as a full replica of its own domain directory partition and the schema and configuration of directory partitions. This global catalog contains a replica of every object in Active Directory, but only a subset of the attributes of each object. In a multiple domain forest with multiple sites, a global catalog server is required for domain logon by anyone but domain administrators.

**infrastructure master**—Responsible for keeping track of updates of group-to-user references, such as a renaming of a user account when group memberships are changed in different domains. The infrastructure master in the group's domain registers the updates and replicates them to other infrastructure masters.

**latency**—The amount of time it takes data to travel from source to destination.

**PDC emulator**—In a mixed-mode domain, the PDC emulator “pretends” to be a Windows NT 4.0 PDC to replicate directory changes to the BDCs in the domain. In a native-mode domain, the PDC emulator also receives preferential replication of password changes from other Windows 2000 domain controllers.

**RID master**—Allocates new relative IDs that are used together with the domain security ID (SID) to create unique security IDs for each object that can be a security principal.

**RMON**—Short for Remote Monitoring, a protocol that allows the monitoring of RMON-enabled hubs and switches from a workstation.

**schema master**—Every domain controller in the forest has a copy of the schema for Active Directory, but only the schema master has a writeable copy. By default, the first domain controller in the first domain in the forest has this role, but this role can be transferred to another domain controller in the forest as needed.

**single master operation roles**—Roles for Active Directory domain controllers.

Roles include the schema master, the domain-naming master, the RID master, the PDC emulator, and the infrastructure master.

---

## REVIEW QUESTIONS

1. You are a consultant with a regional consulting company. You have been asked to mentor a consultant trainee in your organization. Select the statements below that would help to prepare the trainee for network design:
  - a. Most organizations hire our company to completely redesign and replace their existing network, starting from the ground up.
  - b. The network infrastructure is often in place before our consulting company is hired to design a network for new applications and services.
  - c. Business requirements are not important in a network design.
  - d. Company size alone is not very significant in the technical analysis.
  - e. There is no reason to optimize the existing network when you can simply add more bandwidth.
  - f. When you have a T1 connection, you must have performance information to determine what the actual circuit loss is on that connection.
2. Describe the formula for calculating net available bandwidth.
3. Define performance.
4. Define scalability.
5. Define availability.
6. On which leg of the Iron Triangle does scalability belong and why?

Questions 7 and 8 are based on the following scenario:

You are on the design team for a large casino corporation, headquartered in Reno, NV, with casinos in Las Vegas and Reno, NV, and in Kansas City and Lake of the Ozarks, MO. Each casino has a hotel, a nightclub, and a large gift shop. The casinos are open 24 hours a day, and customers at each site can use public kiosks to access reservation information for hotels and entertainment. The company maintains a customer database that includes account information and a history of each customer's expenditures in the casino properties. This information is available to most employees who interact with the customers through the 150 computers that run a third-party client application for the customer database that is stored in a SQL database on a server cluster in the corporate data center in Reno.

7. How would you characterize the security requirements for an organization like this?
8. What are some key considerations for disaster recovery for the casino?

Questions 9 through 12 are based on the following scenario:

The XYZ Corporation, a small manufacturing company with several sites, has hired you as their first full-time network manager. The sites include the corporate headquarters in Yardley, PA, a manufacturing facility in Pennsauken, NJ, and a distribution warehouse in Philadelphia, PA. The company has 400 employees, only 250 of whom need network access. The owner's brother-in-law, a high school math teacher, installed the present network as a part-time job during his summer vacations. It started with a 10 Mbps Ethernet LAN at each location, with ISDN connections to headquarters.

Last summer he converted to 100 Mbps Ethernet, replacing all the hubs, routers, and network cards. There are 10 network printing devices remaining on the network. They use 10 Mbps network cards, but are connected to 10/100 switching hubs.

In spite of the upgrade, the company still experiences some performance problems on the network. They want to add a large new application that will integrate their system for ordering raw materials needed for the manufacturing process with their inventory and manufacturing systems.

9. How would you go about finding the cause of their network performance problems?
10. What was missing in the upgrade of the network?
11. The client does not have an inventory of present applications. What can you do to determine what applications are presently on the network?
12. You have done the analysis for the manufacturing company and determined the applications and services that need to be added to the network as part of the design. How will you determine what load these applications will place on the network?
13. What are the four tasks of capacity planning?
14. Select all of the below that describe network services:
  - a. DHCP
  - b. authentication
  - c. order entry
  - d. accounting
  - e. name resolution
  - f. file and print sharing
15. NetBIOS name resolution is the preferred name resolution in a Windows 2000 domain. True or False? Why?
16. NetBIOS name resolution is not supported in a Windows 2000 domain. True or False? Why?

17. You have optimized a network by eliminating all unnecessary applications and protocols and tracking down and eliminating all unnecessary sources of broadcasts. You have completed load testing of the existing network. Finally, you have done load testing on the additional applications and services your client has requested that you add to the network. The problem is, there is not enough bandwidth available with the current network resources to accommodate all the applications and services, and there is no money in the budget for another four months. What do you do now?
18. What are the possible risks of not having a good centrally planned strategy for IP addresses?
19. You are part of a team working on a network design project for an international company with 50,000 users spread over 130 locations. They currently have Windows NT domains and all Windows clients on the desktops. You are part of the group that must gather hardware and software inventory as part of the technical analysis of the current network. You are told that in preparation for this project, the client installed SMS sites and all computers in the enterprise are SMS clients. You are not savvy about SMS, but you do know what the features are. What features of SMS can you use for your task?
20. Define the difference between latency and bandwidth.

---

## HANDS-ON PROJECTS



### Project 3-1 Calculating Bandwidth

Maple Leaf Candy is a company that provides candy to vending machine suppliers. Their corporate office has a T1 line to a Frame Relay network. Using the formula that is provided in this chapter, perform the appropriate calculations below:

1. Reports from Sprint, your provider of this service, have shown that the average circuit loss over the point-to-point T1 WAN connection during the last six months has been .015 Mbps. Use the bandwidth calculation formula provided here to determine the gross available bandwidth for this WAN link:  
$$1.544 \text{ Mbps} - \text{circuit loss} = \text{Mbps gross available bandwidth}$$
2. Your tests have revealed that the traffic on this link during business hours is .800 Mbps. Although the average is less, use this figure, with the results from your last calculation, to compute the net available bandwidth with the following formula:  
$$\text{gross available bandwidth} - .800 \text{ Mbps} = \text{net available bandwidth}$$
3. Your load tests of the new applications, services, and protocols that will use this link have shown that they require .400 Mbps for sustained periods of time during business hours. Calculate your available bandwidth, using the result from your last calculation in net available bandwidth and the following formula:  
$$\text{net available bandwidth} - 400 \text{ Mbps} = \text{bandwidth cushion available for future needs}$$



## Project 3-2 Testing Network Latency

A simple tool for testing network latency is the ping command. In this lab you will ping another computer from your computer. This project will be more interesting if you can ping a computer that is not on your network subnet, but across a router. If this is possible in your lab, your instructor will give you an address to use. Use this address in place of *remoteaddress* in the following steps and use the IP address of another computer in the lab in place of *localaddress*.

1. Click the **Start** button on the taskbar, and then click **Run**.
2. In the Run box, type **cmd**. The command prompt displays.
3. At the command prompt, type **ping remoteaddress**.
4. Write down the number that indicated the greatest latency.
5. At the command prompt, type **ping localaddress**.
6. Write down the number that indicated the greatest latency.
7. From the command prompt, type **ping /?**. The usage information for ping displays.
8. Notice the various options. Time permitting, experiment with these options.
9. Close all open windows.

These results are more interesting in a production network. This project was intended to teach that a tool we most often use to determine if there is a connection could also be used to display latency information.



## Project 3-3 Using System Monitor for Real-Time Testing

In this lab you will use System Monitor to monitor network performance. Before you begin, you will need the IP address of another student computer in the lab. You will also need to identify a large file or group of files to create shares on your computer and on the other student's computer to use in a file copy. This lab's instructions assume that the shares are already in place.

1. Click the **Start** button on the taskbar.
2. Point to **Settings**.
3. Click **Control Panel**.
4. Double-click **Administrative Tools**.
5. Double-click **Performance**.
6. Right-click the right pane (details pane). The context menu appears.
7. Click **Add Counters**.
8. In the **Performance object** list box, select **Network Interface**. Notice the list of interfaces in the list box on the right. This should include one instance for each NIC in your computer, plus the MS TCP Loopback interface.
9. Be sure that **Select instances from list** is selected, and then click the instance that describes your NIC.

10. On the left, click the **Select counters from list** option button (this is the default).
11. In the list box of performance counters for Network Interface, select **current bandwidth**.
12. Click the **Explain** button at the upper-right of the dialog box and read the explanation of the counter.
13. Click the **Add** button in the upper-right of the dialog box.
14. In the list box of performance counters for Network Interface, select **Packets Outbound Errors** and read the explanation for this counter.
15. Click the **Add** button in the upper-right of the dialog box.
16. In the list box of performance counters for Network Interface, select **Packets Received Non-Unicast/sec** and read the explanation of this counter.
17. Click the **Add** button in the upper-right of the dialog box.
18. In the list box of performance counters for Network Interface, select **Packets Sent Non-Unicast/sec** and read the explanation of this counter.
19. Click the **Add** button in the upper-right of the dialog box.
20. Click **Close**, but leave the windows open on the desktop.
21. Generate traffic between your computer and that of another student by pinging the other computer by IP address (example: ping 192.168.1.200).
22. Generate traffic between your computer and that of another student by copying a large file from your computer to their computer.
23. Generate traffic between your computer and that of another student by copying a large file from the other computer to your computer.
24. View the graph.
25. Close all open windows.



In a classroom environment, you may not get the useful and interesting results you would get in a production environment or a more elaborate test lab.



## Project 3-4 Using System Monitor for Creating Performance Logs

Using System Monitor to view real-time performance data is not as valuable as creating performance logs to establish performance baselines and keep them for comparison when network performance becomes problematic. Performance logs can help when you are doing capacity planning for a network design.

1. Click the **Start** button on the taskbar.
2. Point to **Settings**.

3. Click **Control Panel**.
4. Double-click **Administrative Tools**.
5. Double-click **Performance**.
6. In the Tree pane of the Performance console, click **Performance Logs and Alerts**, and then double-click **Counter logs** in the right pane (the contents pane). System Overview will display. This is a sample settings file for a counter log.
7. Double-click **System Overview**. Although we will not be using this setting file, it is educational to look at the properties to get familiar with what may go into a Counter Log settings file.
8. On the General page, notice the counters that will be logged and the Interval and the Units of the Interval at which data will be sampled.
9. Click the **Log Files** tab.
10. Click the **General** tab again. The first time you do this, a message box displays titled System Overview and complains that the folder c:\PerfLogs does not exist. Click the **Yes** button to allow it to create this directory.
11. Click the **Log Files** tab and notice that the location for the log files will be in the C:\PerLogs directory. Notice the File name and the End file names with option, which will allow you to choose a convention for adding suffixes to the file name.
12. Select the **End file names with:** check box and choose the **mmddhh** format.
13. Notice the example in the box below the choices.
14. Select another format and watch the example change.
15. Select different log file types and you will see the example change to conform with the log file type you select.
16. Click the **Schedule** tab. Notice that you can start logging manually or schedule it to start at a specified date and time.
17. Notice the options for stopping the logging process.
18. Because modifications are not allowed for default logs and alerts, click the **Cancel** button and you can proceed with the following steps in which you will create new log settings to measure network performance.
19. In the Tree pane, right-click **Counter Logs** and select **New Log Settings**.
20. In the Name box of the New Log Settings dialog box, type **Network Baseline**, and then click the **OK** button.
21. The Network Baseline dialog displays. On the General tab, you must add at least one counter before you will be allowed to open another tab sheet in this dialog box.
22. Click the **Add** button. The Select Counters dialog box will display.
23. In the Performance object list box, select **Network Interface**.

24. Verify that **Select Counters from** list is selected, select **Bytes Total/sec**, and ensure that on the right your NIC is selected and that the MS TCP Loopback interface is *not* selected.
25. Click the **Add** button. The Select Counters dialog box will remain open on your screen.
26. In the **Select Counters from** list box, select **Packets/sec** (you will need to use the scroll button to find this counter).
27. Click the **Add** button. The Select Counters dialog box will remain open on your screen.
28. In the Performance object text box, select **Server**.
29. In the **Select Counters from** list box, select **Bytes Total/sec**.
30. Click the **Add** button. The Select Counters dialog box will remain open on your screen.
31. Click the **Close** button to close the Select Counters dialog box.
32. Click the **Log Files** tab.
33. Notice that the location will be C:\PerfLogs. You will not receive an error, because this directory was created in an earlier step. Notice the filename will be Network\_Baseline.
34. If it is not already selected, click the **End file names with** check box, select any suffix you prefer, and allow the numbering to start with 1.
35. In the Log File Type box, select **Text File – CSV** to create a comma separated log file that can be imported into a spreadsheet program for viewing and reporting.
36. Under Log file size, select **Limit of** and keep the default size of 1000 KB.
37. Click the **Schedule** tab.
38. Under the Start log, select the **At** Option Button, and leave the time and date at the present.
39. Under the Stop log, click the **At** option button and configure the stop time to 15 minutes after the start time.
40. Click the **OK** button.
41. Notice that Network Baseline is green, indicating that logging has started.
42. Generate network traffic by pinging other computers and copying the files from your computer to another computer in the lab.
43. After 15 minutes, verify that logging has stopped. The Network Baseline counter log will be red when logging has stopped. If you have a spreadsheet program on your computer, load the resulting data.
44. Close all open windows.



In a classroom environment, you may not get the useful and interesting results you would get in a production environment or a more elaborate test lab.



## Project 3-5 Viewing Network Bindings

Sometimes, as discussed in this chapter, it is necessary to adjust the network binding order. The following steps open the Network Bindings dialog box so that you can view the current bindings.

3

1. Log on to your student machine as administrator.
2. Right-click **My Network Places**. The context menu for My Network Places displays
3. Click **Properties**. The Network and Dial-up Connections dialog box displays.
4. Click the **Advanced** menu on the menu bar at the top of the dialog box.
5. Click **Advanced Settings**.
6. Under Connections, you will see one Local Area Connection for each network interface card in the computer. You will also see Remote Access Connections, even if there is no remote access connection.
7. Click **Local Area Connection** in the Connections list.
8. View the list of bindings under Bindings for Local Area Connection.  
At a minimum, you should see Internet Protocol (TCP/IP) bound to both File and Print Sharing for Microsoft Networks and Client for Microsoft Networks.
9. If there is more than one protocol bound to these components, arrows will appear to the right of the bindings list when one of the protocols is selected. If that is the case, experiment with moving them up and down in the binding order, using the up and down arrows to the right of the bindings list.
10. Make sure that Internet Protocol (TCP/IP) is bound first to both File and Print Sharing for Microsoft Networks and Client for Microsoft Networks before doing the next step.
11. In the Advanced Settings dialog box, click the **Provider Order** tab.
12. Under Network Providers, you should see Microsoft Windows Network and under Print Providers you should see LanMan Print Services and HTTP Print Services.
13. Click **HTTP Print Services**. One of the arrows to the right will become bold.
14. Click that arrow to change the order of HTTP Print Services.
15. Click **OK** to close the Advanced Settings dialog box and accept the changes.
16. Close the **Network and Dial-up Connections** dialog box.
17. Close all open windows.



## Project 3-6 Installing and Using Network Monitor on a Windows 2000 Server

When gathering information for a network design, you may need to look at the traffic on the network. You can use tools such as Sniffer from Sniffer Technologies and Network Monitor that comes with Windows 2000. The following are instructions for installing

Network Monitor and doing a simple network capture. Before you begin, you will need the IP address of the instructor's computer for use in a step near the end of this project.

1. Log on to your Windows 2000 server student computer.
2. Click the **Start** button on the taskbar.
3. Point to **Settings**.
4. Click **Control Panel**.
5. Click **Add/Remove Programs**.
6. Click **Add/Remove Windows Components**.
7. Use the scroll buttons on the right to scroll down until you see Management and Monitoring Tools in the Components list box.
8. Click the words (not the check box) **Management and Monitoring Tools** in the Windows Components Wizard.
9. Click the **Details** button on the lower-right of the Windows Components Wizard.
10. Select the **Network Monitor Tools** check box (be sure not to select any others), and then click **OK**.
11. When prompted for additional files, give the path to the Windows 2000 source files that your instructor will provide, and then click **OK**.
12. When the installation is completed, click the **Finish** button.
13. Click the **Close** button to close Add/Remove Programs.

Both the Network Monitor driver and the Network Monitor management console are installed by the above steps. The console is available on the Administrative Tools menu.

14. To run Network Monitor, click the **Start** button on the taskbar.
15. Point to **Programs**.
16. Point to **Administrative Tools**.
17. Click **Network Monitor**.
18. The first time you run the Network Monitor console, you are prompted to select the default network to monitor.
19. Click **OK** in the Network Monitor – Select Default Network dialog box.
20. In the Select a Network dialog box, click **Local Computer**.
21. If there is only one choice under Local Computer, click that choice, and then skip to Step 31; otherwise, leave the Select a Network dialog box open, and proceed with the next step.
22. Click the **Start** button on the taskbar.
23. Click **Run**.
24. In the Run dialog box, type **cmd**.
25. Press **Enter**.

26. At the command prompt window, type **ipconfig /all**.
27. Write down the value of the physical address under the network adapter you want to monitor (for example: 00-10-4b-94-95-d2): \_\_\_\_\_.
28. Switch back to Network Monitor by clicking **Microsoft Network Monitor** in the taskbar.
29. Under Local Computer, select the network with the value that matches the physical address you recorded above.
30. Click **OK**.
31. The Microsoft Network Monitor window displays with a capture window for your network.
32. Press **F10** to start a capture.
33. Generate some network traffic. If you still have the command prompt open, switch to it. If you do not have the command prompt open, perform Steps 22 through 25. Then continue with Step 35.
34. At the command prompt, type **ping ip address** of the instructor's machine.
35. Copy a file from the instructor's machine to your machine.
36. Switch back to Microsoft Network Monitor.
37. Press **F11** to stop the capture.
38. Close all open windows.

This hands-on project was just an introduction to Network Monitor. For more information on using Network Monitor, check out the Windows 2000 Server Help and [www.microsoft.com/technet](http://www.microsoft.com/technet).

---

## CASE PROJECTS



### Case 3-1 Planning for the Use of Performance Monitor

On your Windows 2000 computer in the classroom or lab, log on as an administrator, open Performance Monitor, and look at the Performance objects. Using the Explain option and the Help program, find several objects and counters you would use to monitor network performance. Find several objects and counters you would use to monitor system performance on a file server.

1. Write a few paragraphs explaining the reasons for using the object and counters you have selected for network performance.
2. Write a few paragraphs explaining the reason for using the objects and counters you have selected for system performance on a file server.



## Case 3-2 Internet Research of Case Studies

Examine case studies at an Internet site and select one to analyze for technical environment and goals, network support for client computer access, and disaster recovery. Write several paragraphs defining the goals and how they were reached in the design. If you feel they were not entirely fulfilled by the design, state your reasons for that conclusion.

Case studies can be found at [www.microsoft.com](http://www.microsoft.com), [www.techrepublic.com](http://www.techrepublic.com) (you might have to join, but it is free), and [www.3com.com](http://www.3com.com).



## Case 3-3 Disaster Recovery

Assume that you work for an insurance company. The insurance company wants to get into the business of writing policies that cover business losses associated with disasters that damage computer equipment. In lay terms, explain to the insurance company's management the difference between client computers and servers and why a company without a disaster recovery plan is a high-risk client. Establish a checklist that the insurance company can use to determine if a client is a high-risk client or a low-risk client.